

JCE20 U.S. PTO
09/29/00

10-02-00

A

Practitioner's Docket No. 55,045 (158)

PATENT

JCE20 U.S. PTO
09/29/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of
Inventor(s): Brian James Martin and Michael Ryan

WARNING: 37 C.F.R. § 1.41(a)(1) points out:

"(a) A patent is applied for in the name or names of the actual inventor or inventors.

(1) The inventorship of a nonprovisional application is that inventorship set forth in the oath or declaration as prescribed by § 1.63, except as provided for in § 1.53(d)(4) and § 1.63(d). If an oath or declaration as prescribed by § 1.63 is not filed during the pendency of a nonprovisional application, the inventorship is that inventorship set forth in the application papers filed pursuant to § 1.53(b), unless a petition under this paragraph accompanied by the fee set forth in § 1.17(i) is filed supplying or changing the name or names of the inventor or inventors."

For (title): METHOD OF REDUCING CONTENTION OF A HIGHLY CONTESTED LOCK
PROTECTING MULTIPLE DATA ITEMS

1. Type of Application

This new application is for a(n)

CERTIFICATION UNDER 37 C.F.R. 1.10*

(Express Mail label number is mandatory.)

(Express Mail certification is optional.)

I hereby certify that this correspondence and the documents referred to as attached therein are being deposited with the United States Postal Service on this date, September 29, 2000, in an envelope as "Express Mail Post Office to Addressee," mailing Label Number EL180585765US, addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

William J. Daley, Jr.
(type or print name of person mailing paper)

William J. Daley
Signature of person mailing paper

WARNING: Certificate of mailing (first class) or facsimile transmission procedures of 37 C.F.R. 1.8 cannot be used to obtain a date of mailing or transmission for this correspondence.

***WARNING:** Each paper or fee filed by "Express Mail" must have the number of the "Express Mail" mailing label placed thereon prior to mailing. 37 C.F.R. 1.10(b).
"Since the filing of correspondence under § 1.10 without the Express Mail mailing label thereon is an oversight that can be avoided by the exercise of reasonable care, requests for waiver of this requirement will not be granted on petition." Notice of Oct. 24, 1996, 60 Fed. Reg. 56,439, at 56,442.

(check one applicable item below)

- [X] Original (nonprovisional)
[] Design
[] Plant

WARNING: *Do not use this transmittal for a completion in the U.S. of an International Application under 35 U.S.C. 371(c)(4), unless the International Application is being filed as a divisional, continuation or continuation-in-part application.*

WARNING: *Do not use this transmittal for the filing of a provisional application.*

NOTE: *If one of the following 3 items apply, then complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF A PRIOR U.S. APPLICATION CLAIMED and a NOTIFICATION IN PARENT APPLICATION OF THE FILING OF THIS CONTINUATION APPLICATION.*

- [] Divisional.
[] Continuation.
[] Continuation-in-part (C-I-P).

2. Benefit of Prior U.S. Application(s) (35 U.S.C. 119(e), 120, or 121)

NOTE: *A nonprovisional application may claim an invention disclosed in one or more prior filed copending nonprovisional applications or copending international applications designating the United States of America. In order for a nonprovisional application to claim the benefit of a prior filed copending nonprovisional application or copending international application designating the United States of America, each prior application must name as an inventor at least one inventor named in the later filed nonprovisional application and disclose the named inventor's invention claimed in at least one claim of the later filed nonprovisional application in the manner provided by the first paragraph of 35 U.S.C. 112. Each prior application must also be:*

- (i) *An international application entitled to a filing date in accordance with PCT Article 11 and designating the United States of America; or*
- (ii) *Complete as set forth in § 1.51(b); or*
- (iii) *Entitled to a filing date as set forth in § 1.53(b) or § 1.53(d) and include the basic filing fee set forth in § 1.16; or*
- (iv) *Entitled to a filing date as set forth in § 1.53(b) and have paid therein the processing and retention fee set forth in § 1.21(l) within the time period set forth in § 1.53(f).*

37 C.F.R. § 1.78(a)(1).

NOTE *If the new application being transmitted is a divisional, continuation or a continuation-in-part of a parent case, or where the parent case is an International Application which designated the U.S., or benefit of a prior provisional application is claimed, then check the following item and complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.*

WARNING: *If an application claims the benefit of the filing date of an earlier filed application under 35 U.S.C. 120, 121 or 365(c), the 20-year term of that application will be based upon the filing date of the earliest U.S. application that the application makes reference to under 35 U.S.C. 120, 121 or 365(c). (35 U.S.C. 154(a)(2) does not take into account, for the determination of the patent term, any application on which priority is claimed under 35 U.S.C. 119, 365(a) or 365(b).) For a c-i-p application, applicant should review whether any claim in the patent that will issue is supported by an earlier application and, if not, the applicant should consider canceling the reference to the earlier filed application. The term of a patent is not based on a claim-by-claim approach. See Notice of April 14, 1995, 60*

Fed. Reg. 20,195, at 20,205.

WARNING: When the last day of pendency of a provisional application falls on a Saturday, Sunday, or Federal holiday within the District of Columbia, any nonprovisional application claiming benefit of the provisional application must be filed prior to the Saturday, Sunday, or Federal holiday within the District of Columbia. See 37 C.F.R. § 1.78(a)(3).

- [] The new application being transmitted claims the benefit of prior U.S. application(s).

3. Papers Enclosed

A. Required for Filing Date under 37 C.F.R. § 1.53(b) (Regular) or 37 C.F.R. § 1.153 (Design) Application

21 Pages of Specification
9 Pages of Claims
7 Sheets of Drawing (informal)

WARNING: *DO NOT submit original drawings. A high quality copy of the drawings should be supplied when filing a patent application. The drawings that are submitted to the Office must be on strong, white, smooth, and non-shiny paper and meet the standards according to § 1.84. If corrections to the drawings are necessary, they should be made to the original drawing and a high-quality copy of the corrected original drawing then submitted to the Office. Only one copy is required or desired. For comments on proposed then-new 37 C.F.R. 1.84, see Notice of March 9, 1988. (1990 O.G. 57-62).*

NOTE: "Identifying indicia, if provided, should include the application number or the title of the invention, inventor's name, docket number (if any), and the name and telephone number of a person to call if the Office is unable to match the drawings to the proper application. This information should be placed on the back of each sheet of drawing a minimum distance of 1.5 cm. (5/8 inch) down from the top of the page..." 37 C.F.R. § 1.84(c)).

(complete the following, if applicable)

- [] The enclosed drawing(s) are photograph(s), and there is also attached a "PETITION TO ACCEPT PHOTOGRAPH(S) AS DRAWING(S)." 37 C.F.R. § 1.84(b).

- [] Formal
[X] Informal

B. Other Papers Enclosed

1 Pages of Abstract
_____ Other

4. Additional Papers Enclosed

[] Amendment to claims

- [] Cancel in this application claims _____ before calculating the filing fee. (At least one original independent claim must be retained for filing purposes.)
- [] Add the claims shown on the attached amendment. (Claims added have been numbered consecutively following the highest numbered original claims.)

[] Preliminary Amendment

[] Information Disclosure Statement (37 C.F.R. § 1.98)

[] Form PTO-1449 (PTO/SB/08A and 08B)

[] Citations

[] Declaration of Biological Deposit

[] Submission of "Sequence Listing," computer readable copy and/or amendment pertaining thereto for biotechnology invention containing nucleotide and/or amino acid sequence.

[] Authorization of Attorney(s) to Accept and Follow Instructions from Representative

[] Special Comments

[] Other

5. Declaration or Oath (including power of attorney)

NOTE: A newly executed declaration is not required in a continuation or divisional application provided the prior nonprovisional application contained a declaration as required, the application being filed is by all or fewer than all the inventors named in the prior application, there is no new matter in the application being filed, and a copy of the executed declaration filed in the prior application (showing the signature or an indication thereon that it was signed) is submitted. The copy must be accompanied by a statement requesting deletion of the names of person(s) who are not inventors of the application being filed. If the declaration in the prior application was filed under § 1.47 then a copy of that declaration must be filed accompanied by a copy of the decision granting § 1.47 status or, if a nonsigning person under § 1.47 has subsequently joined in a prior application, then a copy of the subsequently executed declaration must be filed. See 37 C.F.R. § 1.63(d)(1)-(3).

NOTE: A declaration filed to complete an application must be executed, identify the specification to which it is directed, identify each inventor by full name, including the family name, and at least one given name without abbreviation together with any other given name or initial, and the residence, post office address and country of citizenship of each inventor, and state whether the inventor is a sole or joint inventor. 37 C.F.R. § 1.63(a)(1)-(4).

[X] Enclosed (unexecuted).

Executed by

(check all applicable boxes)

- [] inventor(s).
- [] legal representative of inventor(s). 37 C.F.R. § 1.42 or 1.43.
- [] joint inventor or person showing a proprietary interest on behalf of inventor who refused to sign or cannot be reached.

This is the petition required by 37 C.F.R. § 1.47 and the statement required by 37 C.F.R. § 1.47 is also attached. See item 13 below for fee.

Not Enclosed.

NOTE: *Where the filing is a completion in the U.S. of an International Application, or where the completion of the U.S. application contains subject matter in addition to the International Application, the application may be treated as a continuation or continuation-in-part, as the case may be, utilizing ADDED PAGE FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION CLAIMED.*

Application is made by a person authorized under 37 C.F.R. 1.41(c) on behalf of *all* the above named inventor(s).

(The declaration or oath, along with the surcharge required by 37 C.F.R. § 1.16(e), can be filed subsequently).

Showing that the filing is authorized.
(not required unless called into question. 37 C.F.R. § 1.41(d))

6. Inventorship Statement

WARNING: *If the named inventors are each not the inventors of all the claims an explanation, including the ownership of the various claims at the time the last claimed invention was made, should be submitted.*

The inventorship for all the claims in this application are:

- The same.
or
 Not the same. An explanation, including the ownership of the various claims at the time the last claimed invention was made,
 is submitted.
 will be submitted.

7. Language

NOTE: *An application including a signed oath or declaration may be filed in a language other than English. An English translation of the non-English language application and the processing fee of \$130.00 required by 37 C.F.R. § 1.17(k) is required to be filed with the application, or within such time as may be set by the Office. 37 C.F.R. § 1.52(d).*

- English
 Non-English

The attached translation includes a statement that the translation is accurate. 37 C.F.R. § 1.52(d).

8. Assignment

[X] An assignment of the invention to EMC Corporation was recorded on _____ reel _____, frame _____ (copy enclosed).

[] is attached. A separate [] "COVER SHEET FOR ASSIGNMENT (DOCUMENT) ACCOMPANYING NEW PATENT APPLICATION" or [] FORM PTO 1595 is also attached.

[X] will follow.

NOTE: "If an assignment is submitted with a new application, send two separate letters-one for the application and one for the assignment" Notice of May 4, 1990 (1114 O.G. 77-78).

WARNING: A newly executed "STATEMENT UNDER 37 C.F.R. § 3.73(b)" must be filed when a continuation-in-part application is filed by an assignee. Notice of April 30, 1993, 1150 O.G. 62-64.

9. Certified Copy

Certified copy(ies) of application(s)

Country	Appln. no.	Filed

from which priority is claimed

[] is (are) attached.

[] will follow.

[] were filed on _____, in prior application U.S.S.N. _____, which was filed on _____.

NOTE: The foreign application forming the basis for the claim for priority must be referred to in the oath or declaration. 37 C.F.R. § 1.55(a) and 1.63.

NOTE: This item is for any foreign priority for which the application being filed directly relates. If any parent U.S. application or International Application from which this application claims benefit under 35 U.S.C. 120 is itself entitled to priority from a prior foreign application, then complete item 18 on the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

10. Fee Calculation (37 C.F.R. § 1.16)

A. Regular application

CLAIMS AS FILED

Claims	Number Filed	Basic Fee Allowance	Number Extra	Rate	Basic Fee 37 C.F.R. § 1.16(a) \$690.00
Total Claims (37 C.F.R. § 1.16(c))	29	- 20 =	9 x	\$ 18.00	
					\$162.00
Independent Claims (37 C.F.R. § 1.16(b))	5	- 3 =	2 x	\$ 78.00	
					\$156.00
Multiple Dependent Claim(s), if any (37 C.F.R. § 1.16(d))			0 +	\$260.00	

- Amendment cancelling extra claims is enclosed.
 Amendment deleting multiple-dependencies is enclosed.
 Fee for extra claims is not being paid at this time.

NOTE: If the fees for extra claims are not paid on filing they must be paid or the claims cancelled by amendment, prior to the expiration of the time period set for response by the Patent and Trademark Office in any notice of fee deficiency. 37 C.F.R. § 1.16(d).

Filing Fee Calculation \$ 1008.00

B. Design application
(\$310.00—37 C.F.R. § 1.16(f))

Filing Fee Calculation \$ _____

C. Plant application
(\$480.00—37 C.F.R. § 1.16(g))

Filing Fee Calculation \$ _____

11. Small Entity Statement(s)

- Statement(s) that this is a filing by a small entity under 37 C.F.R. §§ 1.9 and 1.27 will be submitted at a later date.

WARNING: *"Status as a small entity must be specifically established in each application or patent in which the status is available and desired. Status as a small entity in one application or patent does not affect any other application or patent, including applications or patents which are directly or indirectly dependent upon the application or patent in which the status has been established. The refiling of an application under § 1.53 as a continuation, division, or continuation-in-part (including a continued prosecution application under § 1.53(d)), or the filing of a reissue application requires a new determination as to continued entitlement to small entity status for the continuing or reissue application. A nonprovisional application claiming benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) of a prior application, or a reissue application may rely on a statement filed in the prior application or in the patent if the nonprovisional application or the reissue application includes a reference to the statement in the prior application or in the patent or includes a copy of the statement in the prior application or in the patent and status as a small entity is still proper and desired. The payment of the small entity basic statutory filing fee will be treated as such a reference for purposes of this section." 37 C.F.R. § 1.28(a)(2).*

(complete the following, if applicable)

- Status as a small entity was claimed in prior application _____ / _____, filed on _____ from which benefit is being claimed for this application under:

35 U.S.C. § 119(e),
 120,
 121,
 365(c),

and which status as a small entity is still proper and desired.

- A copy of the statement in the prior application is included.

Filing Fee Calculation (50% of **A**, **B** or **C** above) \$_____

NOTE: *Any excess of the full fee paid will be refunded if a small entity status is established refund request are filed within 2 months of the date of timely payment of a full fee. The two-month period is not extendable under § 1.136. 37 C.F.R. § 1.28(a).*

12. Request for International-Type Search (37 C.F.R. § 1.104(d))

(complete, if applicable)

- Please prepare an international-type search report for this application at the time when national examination on the merits takes place.

13. Fee Payment Being Made at This Time

Not Enclosed

No filing fee is to be paid at this time.

(This and the surcharge required by 37 C.F.R. § 1.16(e) can be paid subsequently.)

Enclosed

Filing fee \$ _____

Recording assignment
(\$40.00; 37 C.F.R. § 1.21(h))
(See attached "COVER SHEET FOR
ASSIGNMENT ACCOMPANYING NEW
APPLICATION.") \$ _____

Petition fee for filing by other
than all the inventors or person
on behalf of the inventor where
inventor refused to sign or cannot
be reached
(\$130.00; 37 C.F.R. §§ 1.47 and 1.17(i)) \$ _____

For processing an application with a
specification in a non-English language
(\$130.00; 37 C.F.R. §§ 1.52(d) and 1.17(k)) \$ _____

Processing and retention fee
(\$130.00; 37 C.F.R. §§ 1.53(d) and 1.21(l)) \$ _____

Fee for international-type search report
(\$40.00; 37 C.F.R. § 1.21(e)) \$ _____

NOTE: 37 C.F.R. § 1.21(l) establishes a fee for processing and retaining any application that is abandoned for failing to complete the application pursuant to 37 C.F.R. § 1.53(f) and this, as well as the changes to 37 C.F.R. § 1.53 and 1.78(a)(1), indicate that in order to obtain the benefit of a prior U.S. application, either the basic filing fee must be paid, or the processing and retention fee of § 1.21(l) must be paid, within 1 year from notification under § 53(f).

Total Fees Enclosed \$ _____

14. Method of Payment of Fees

- [] Check in the amount of \$_____.
- [] Charge Account No. _____ in the amount of \$_____.
A duplicate of this transmittal is attached.

NOTE: Fees should be itemized in such a manner that it is clear for which purpose the fees are paid. 37 C.F.R. § 1.22(b).

15. Authorization to Charge Additional Fees

WARNING: *If no fees are to be paid on filing, the following items should not be completed.*

WARNING: *Accurately count claims, especially multiple dependent claims, to avoid unexpected high charges, if extra claim charges are authorized.*

- [] The Commissioner is hereby authorized to charge the following additional fees by this paper and during the entire pendency of this application to Account No._____.
- [] 37 C.F.R. § 1.16(a), (f) or (g) (filing fees)
- [] 37 C.F.R. § 1.16(b), (c) and (d) (presentation of extra claims)

NOTE: Because additional fees for excess or multiple dependent claims not paid on filing or on later presentation must only be paid or these claims cancelled by amendment prior to the expiration of the time period set for response by the PTO in any notice of fee deficiency (37 C.F.R. § 1.16(d)), it might be best not to authorize the PTO to charge additional claim fees, except possibly when dealing with amendments after final action.

- [] 37 C.F.R. § 1.16(e) (surcharge for filing the basic filing fee and/or declaration on a date later than the filing date of the application)
- [] 37 C.F.R. § 1.17(a)(1)-(5) (extension fees pursuant to § 1.136(a)).
- [] 37 C.F.R. § 1.17 (application processing fees)

NOTE: "A written request may be submitted in an application that is an authorization to treat any concurrent or future reply, requiring a petition for an extension of time under this paragraph for its timely submission, as incorporating a petition for extension of time for the appropriate length of time. An authorization to charge all required fees, fees under § 1.17, or all required extension of time fees will be treated as a constructive petition for an extension of time in any concurrent or future reply requiring a petition for an extension of time under this paragraph for its timely submission. Submission of the fee set forth in § 1.17(a) will also be treated as a constructive petition for an extension of time in any concurrent reply requiring a petition for an extension of time under this paragraph for its timely submission." 37 C.F.R. § 1.136(a)(3).

- [] 37 C.F.R. § 1.18 (issue fee at or before mailing of Notice of Allowance, pursuant to 37 C.F.R. § 1.311(b))

NOTE: Where an authorization to charge the issue fee to a deposit account has been filed before the mailing of a Notice of Allowance, the issue fee will be automatically charged to the deposit account at the time of mailing the notice of allowance. 37 C.F.R. § 1.311(b)).


NOTE: 37 C.F.R. § 1.28(b) requires "Notification of any change in status resulting in loss of entitlement to small entity status must be filed in the application . . . prior to paying, or at the time of paying, . . . issue fee." From the wording of 37 C.F.R. § 1.28(b), (a) notification of change of status must be made even if the fee is paid as "other than a small entity" and (b) no notification is required if the change is to another small entity.

16. Instructions as to Overpayment

NOTE: ". . . Amounts of twenty-five dollars or less will not be returned unless specifically requested within a reasonable time, nor will the payer be notified of such amounts; amounts over twenty-five dollars may be returned by check or, if requested, by credit to a deposit account." 37 C.F.R. § 1.26(a).

Credit Account No. _____.

Refund


SIGNATURE OF PRACTITIONER

Reg. No. 35,487

Tel. No. (617) 523-3400

Customer No.: 21,874

William J. Daley, Jr.
Dike, Bronstein, Roberts & Cushman, LLP
Intellectual Property Practice Group
EDWARDS & ANGELL, LLP
130 Water Street
Boston, MA 02109

Incorporation by reference of added pages

(check the following item if the application in this transmittal claims the benefit of prior U.S. application(s) (including an international application entering the U.S. stage as a continuation, divisional or C-I-P application) and complete and attach the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED)

- Plus Added Pages for New Application Transmittal Where Benefit of Prior U.S. Application(s) Claimed

Number of pages added _____

- Plus Added Pages for Papers Referred to in Item 4 Above

Number of pages added _____

- Plus added pages deleting names of inventor(s) named on prior application(s) who is/are no longer inventor(s) of the subject matter claimed in this application.

Number of pages added _____

- Plus "Assignment Cover Letter Accompanying New Application"

Number of pages added _____

Statement Where No Further Pages Added

(if no further pages form a part of this Transmittal, then end this Transmittal with this page and check the following item)

- This transmittal ends with this page.

CERTIFICATE OF MAILING BY "EXPRESS MAIL" (37 CFR 1.10)

Applicant(s): Brian James Martin and Michael Ryan

Docket No.

55,045 (158)

Serial No.
Not Yet AssignedFiling Date
HerewithExaminer
Not Yet AssignedGroup Art Unit
Not Yet AssignedInvention: **METHOD OF REDUCING CONTENTION OF A HIGHLY CONTENDED LOCK PROTECTING
MULTIPLE DATA ITEMS**

I hereby certify that this new U.S. utility patent application
(Identify type of correspondence)

is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under
37 CFR 1.10 in an envelope addressed to: The Assistant Commissioner for Patents, Washington, D.C. 20231 on
September 29, 2000.
(Date)

William J. Daley, Jr.*(Typed or Printed Name of Person Mailing Correspondence)*

A handwritten signature of William J. Daley, Jr. is written over a horizontal line.

*(Signature of Person Mailing Correspondence)*EL180585765US*("Express Mail" Mailing Label Number)*

Note: Each paper must have its own certificate of mailing.

Docket No. 55,045 (158)
Express Mail Label No. EL180585765US

METHOD OF REDUCING CONTENTION OF A HIGHLY CONTESTED LOCK
5 PROTECTING MULTIPLE DATA ITEMS

FIELD OF INVENTION

The present invention relates to methods and processes for protecting multiple data item using locks and more particularly to a method/ process for reducing the 10 contention of a lock(s).

BACKGROUND OF THE INVENTION

Many current computer systems employ a multi-processor configuration that includes two or more processing units interconnected by a bus system and each being 15 capable of independent or cooperative operation. Such a multi-processor configuration increases the total system processing capability and allows the concurrent execution of multiple related or separate tasks by assigning each task to one or more processors. Such systems also typically include a plurality of mass storage units, such as disk drive devices to provide adequate storage capacity for the 20 number of task executing on the systems.

One type of multi-processor computer system embodies a symmetric multiprocessing (SMP) computer architecture which is well known in the art as overcoming the limitations of single or uni-processors in terms of processing speed and transaction throughput, among other things. Typical, commercially available 25 SMP systems are generally "shared memory" systems, characterized in that multiple processors on a bus, or a plurality of busses, share a single global memory or shared

memory. In shared memory multiprocessors, all memory is uniformly accessible to each processor, which simplifies the task of dynamic load distribution. Processing of complex tasks can be distributed among various processors in the multiprocessor system while data used in the processing is substantially equally available to each of the processors undertaking any portion of the complex task. Similarly, programmers writing code for typical shared memory SMP systems do not need to be concerned with issues of data partitioning, as each of the processors has access to and shares the same, consistent global memory.

There is shown in FIG. 1 a block diagram of an exemplary multiprocessor system that implements an SMP architecture. For further details regarding this system, reference shall be made to U.S.S.N. 09/309,012, filed September 3, 1999, the teachings of which are incorporated herein by reference.

Another computer architecture known in the art for use in a multi-processor environment is the Non-Uniform Memory Access (NUMA) architecture or the Cache Coherent Non-Uniform Memory Access (CCNUMA) architecture, which are known in the art as being an extension of SMP but which supplants SMPs “shared memory architecture.” NUMA and CCNUMA architectures are typically characterized as having distributed global memory. Generally, NUMA/ CCNUMA machines consist of a number of processing nodes connected through a high bandwidth, low latency interconnection network. The processing nodes are each comprised of one or more high-performance processors, associated cache, and a portion of a global shared memory. Each node or group of processors has near and far memory, near memory being resident on the same physical circuit board, directly accessible to the node's processors through a local bus, and far memory being resident on other nodes and

being accessible over a main system interconnect or backbone. Cache coherence, i.e. the consistency and integrity of shared data stored in multiple caches, is typically maintained by a directory-based, write-invalidate cache coherency protocol, as known in the art. To determine the status of caches, each processing node typically has a
5 directory memory corresponding to its respective portion of the shared physical memory. For each line or discrete addressable block of memory, the directory memory stores an indication of remote nodes that are caching that same line.

There is shown in FIG. 2 a high-level block diagram of another exemplary multiprocessor system but which implements a CCNUMA architecture. For further details regarding this system, reference shall be made to U.S.P. 5,887,146, the teachings of which are incorporated herein by reference.
10
15

The operating systems for such multiprocessor systems or the user application(s) for execution on such multiprocessor systems can employ a methodology whereby a lock is used to protect multiple data items, and/or multiple instances of a data item that are in the memory (e.g., RAM) of the multiprocessor system. For example, a data set that includes the telephone number for a user(s), user name(s) and addresses of the user(s). Each time the operating system or user application accesses the data or data items in memory, such as for a read or a writing operation, a global lock is acquired (STEP 2, FIG. 3) over the data/ data items. After
20 the global lock is acquired, the user application or operating system accesses the data/ data items for example in one of a read/ write operation (STEP 4, FIG. 3). After accessing the data/ data items, the global lock is released (STEP 6, FIG. 3).

The foregoing process shown in FIG. 3 is intended to ensure that the data or data items are not changed during the time period when the data is being accessed,

for example, for a read or write operation. In other words, more than one read/ write operation cannot be performed at the same time. After the global lock is released the operating system or user application can perform this process again for the next or another accessing of data/ data items.

5 Although this techniques is simple, acquiring a global lock each and every time data is to be accessed for a read/ write operation for example, becomes increasingly more time consuming as more concurrent accessing operations are initiated and therefore contend for the lock. As also indicated above, while the accessing of data in the memory for one given accessing operation is being performed, other accessing operations trying to access the same data cannot be performed, thus the applications programs/ operating system/ processors involved with such other accessing operations are unable to proceed (i.e., operation pended or delayed). Further, in the case where the global lock is obtained over a data set comprising multiple data items, the other accessing operations can be for data items not involved with the one given accessing operation being performed.

For example, lets assume that the data being stored in the memory is a phonebook type of listing, that includes name, address and telephone number of all subscribers, and which is accessed by any one of a number of operators to obtain listing information to give to callers. When one operator accesses the data for one subscriber, a global lock is obtained over the data for all subscribers. Similarly, if a person is updating the data provided for a given subscriber, a global lock is obtained over all of the data for all subscribers until the updating operation is completed and the global lock released. Consequently, the next operator attempting to access the

data must await the release of this global lock before the next accessing of the data can proceed.

If one or more applications programs and/or the operating system that are being run on a multiprocessor system require or involve frequent accessing of 5 common data in the memory, then the various data accessing operations can become in competition with each other. As a consequence, the competing data accessing operations can cause the processing of the one or more data operations to be delayed.

Consequently, the time to perform a task by an applications program and/or an operating system is in effect increased. Although the foregoing is described in connection with specific multiprocessor system implementations, it should be recognized that similar time delays can occur in other multiprocessor system configurations in which an applications program or operating system experiences data access global lock contention as described herein.

It thus would be desirable to provide new methodologies or techniques for optimizing applications programs and/or operating systems so as to reduce such global lock contentions while ensuring the data being accessed from the memory is not corrupted or changed during the time period when the data is being accessed. Further, it would be desirable to provide such methods that would reduce the amount of time to perform tasks being performed by the applications program or operating 20 systems as compared to prior art methods and techniques.

SUMMARY OF THE INVENTION

The present invention features a method or process for reducing contention of a highly contended software lock(s) that is protecting multiple data items, where the

software has a plurality of code paths accessing the data items. The method includes identifying one or more heavily used code paths and determining which data items are touched by the identified one or more heavily used code paths. These data items are then moved into a partition, if they were not partitioned earlier.

5 The locking requirements for each of the identified one or more heavily used code paths are optimized to match the reduced locking requirements because of the partitioned data items. In other words, the locking requirements can be reduced so only the locks for the partitions including the data items that touch the code path are acquired. Such a method further includes acquiring all partition locks and the global lock, wherever a global lock would have been acquired to protect data.

10 In more specific embodiments of the present invention, the so-modified software is evaluated to determine if there is an acceptable increase in overall system performance resulting from the optimization of the locking requirements in a selected one or more heavily used code paths. If the system's performance following optimization is not acceptable, then the next most heavily used code path is identified and the locking requirements for this code path are optimized. Such optimization is continued until the system exhibits an acceptable overall system performance or all code paths are optimized.

15 According to another aspect of the present invention, the method or process for reducing contention of a highly contended software lock that is protecting multiple data items is arranged so as to optimize the locking requirements for each of the code paths of the software.

20 Other aspects and embodiments of the invention are discussed below.

BRIEF DESCRIPTION OF THE DRAWING

For a fuller understanding of the nature and desired objects of the present invention, reference is made to the following detailed description taken in conjunction with the accompanying drawing figures wherein like reference character denote corresponding parts throughout the several views and wherein:

FIG. 1 is a block diagram of an exemplary multiprocessor system that implements a SMP architecture; and

FIG. 2 is a high-level block diagram of another exemplary multiprocessor system that implements a CCNUMA architecture

FIG. 3 is a flow diagram of a conventional technique for protecting data stored in memory;

FIG. 4 is a high-level flow diagram illustrating the process or method of the present invention for optimizing an applications program or operating system;

FIGS. 5A-C are illustrative of various states of software being optimized;

FIG. 6 is a flow diagram of the lock-all locking routine; and

FIG. 7 is a high-level flow diagram illustrating another process or method of the present invention for optimizing an applications program or operating system.

DESCRIPTION OF THE PREFERRED EMBODIMENT

As noted above, the software developed for operating systems and/or applications programs for execution in computer systems, particularly multiprocessor computer systems, include locks protecting multiple data items of a data set or multiple instances of a data item. Such locks are typically referred to as global locks that protect one or more sets of data. For purposes of the present invention, the

contents of the set of data or data set can be any of a number of data items known to those skilled in the art. Further, although the following may be seen as describing a methodology for optimizing already written software, it should be recognized that the methodology and techniques of the present invention can be implemented or effected 5 so as to be performed in conjunction with the writing of the code for software.

It also shall be understood that the software comprising applications programs or operating systems includes one or more code paths that access data protected by a global lock. In more specific embodiments, the software includes a plurality or multiplicity or more of such code paths. For purposes of the present invention, a software code path according to the present invention begins at a location in the software code that acquires the global lock, and ends when the global lock is released.

As such, each code path ordinarily employing such conventional locking techniques would involve the locking and unlocking of all items within the data set before and after respectively accessing the locked data much in the fashion illustrated in FIG. 3. Also, the data set can comprise one or more data items that can be accessed across multiple code paths. In more specific embodiments, the data set includes four or more data items.

There is shown in FIG. 4 a high-level flow diagram that illustrates the process or method of the present invention for optimizing the software for an applications program or operating system so as to reduce the contention of a highly contended lock, a global lock, protecting multiple data items. The present method begins by 20 determining the methodology by which the data set is to be partitioned, STEP 402.

As indicated above, when using conventional techniques a programmer(s) or a software designer(s) typically provides a global lock in each code path accessing a data

set in the memory (RAM) of a computer system to protect all of the data contained in the data set. It should be recognized that the data or contents of such a data set do not particularly limit a method according to the present invention.

In accordance with the method of the present invention, a programmer(s) or software designer(s) evaluates any one or more of the computer system and/or the particular operational characteristics, functions and tasks of the software as well as the data of the data set, including the kind and type of data, that is stored in the system memory (RAM, SDRAM, etc.) of the computer system. From this evaluation, the programmer determines how to in effect divide or partition this stored memory into subsets or partitions. In general terms, the number of data subsets or partitions and the definition of the contents of each subset or partition are established so as to provide a mechanism to redefine the locking requirements of one or more code paths of the software. The redefinition of the locking requirements of one or more code paths, further provides a mechanism for reducing contention of a highly contended lock that would have been provided if one used conventional software techniques for accessing data for the one or more code paths. The following provides some exemplary data set/ system configurations and some exemplary partitioning schemes that one skilled in the art might contemplate for use in connection with the methodology of the present invention. Such discussion also further describes the methodology of the present invention.

In the case of a NUMA type of computer system, such a system typically includes a plurality of processing nodes or physical locales and distributed global memory, where each processing node includes a portion of the global memory. Each processing node, or a group of processing nodes, has near and far memory, the near

memory being accessible to the node's processors through a local bus and far memory at other processing nodes and being accessible over a main system bus or other interconnect. Although there is near and far memory for such a system, when applying conventional techniques for protecting data in a data set a global lock would 5 be applied across the entire distributed global memory to protect the data within the data set.

Assume also for example, that the function or tasks being performed at each processing node, involves the use of data that resides in the near memory of each processing node. In such a case, the number of data subsets or partitions and the definition of the contents of each data subset or partition is related to the number of physical locales or processing nodes of the NUMA computer system and the data already found in the near memory of each processing node. In other words, the number of partitions and the contents of each partition are determined based on the physical arrangement and number of the NUMA processing nodes and the way in which data is already distributed and stored in the global memory of the NUMA computer system. Thus, for example, the data in the memory of the first processing node would be partitioned from the data in the memory of the second and subsequent processing nodes.

In the case of a data set comprising an information listing, for example a 20 phonebook type of listing, one of the fields of data may include information that inherently provides a mechanism or way of sub-dividing the data into discrete subsets. For example, in the case of the phonebook-listing example the number and definition of each data subset can be based on the alphabet and the first letter of the last name of the telephone subscribers. Thus, for example, the data comprising the

phonebook listing would be partitioned so the data corresponding to all listings having last names beginning with A would be partitioned from the data corresponding to all listings having last names beginning with B-Z (i.e., the other letters of the alphabet).

5 In addition, the partitioning of the data set can be determined or established based on a consideration of a specific parameter or specific data field that is utilized in the software code for one or more code paths. Referring to the phonebook information listing, a particular applications program can include a number of code paths to perform various tasks or functions, including updating the listing for changes in subscribers (adds/ deletes/ revisions) and reading of the telephone number for a given subscriber name and address. Also assuming that the heaviest code path from a usage standpoint is that involving the reading of the telephone number, then the data in the data subset would be partitioned so as to be capable of using the data field that would have the most impact on reducing contention. As with the above, the number and definition of the data subsets or partitions can be based on the alphabet and the first letter of the last name of the telephone subscribers.

10 In the case of an operating system, suppose there exists a system for managing a range of memory, where each page in the range is described by a structure that contains some information on the status and/or contents of that page. Using a conventional technique for protecting data, a single global lock would be utilized to protect the many instances of the page descriptor structure. In the present method, N partitions would be created and the instances of the page descriptor structure would be divided among the N partitions. In this way, a partition is defined as a range of memory and the items associated with that memory range.

In sum, such partitions are not particularly limited to a particular form.

Further such partitions, in exemplary embodiments, can be set so as to correspond or relate to a set of processors of a multiprocessor computer system, correspond or relate to a physical memory range, correspond or relate to a specific existing relationship amongst the data within the data set or be entirely conceptual in nature with no physical basis. The concept of a partition in the present invention refers to the set of data locked only by a particular partition lock, and specifically does not imply any rearrangement or relocation of the data itself.

After determining the partitioning method, N partition locks are established or created in the software for one or more data sets in the system memory (RAM) to be accessed in the one or more code paths of the software, step 404, where $N \geq 2$. Such partition locks are created with the same locking mechanism employed by the original global lock or using any other technique known to those skilled in the art that is otherwise compatible with the software code being modified or written. There is shown in FIG. 5A, the state of the software with the data set protected by the global lock and the newly created, empty partitions and the associated partition locks.

Following the creation of the N partition locks, the software is modified or is written such that wherever the global lock was acquired, all N partition locks and the global lock are now acquired, Step 406. With this provision in the software, the data comprising the original, or initial, data set is effectively defined so as to be protected by both the global lock and all of the partition locks. The state of the software with the data set protected by the global lock and the newly created, empty partitions and the associated partition locks is illustrated in FIG. 5B.

10
15
20

Because a plurality of locks may be held at the same time in any one or more code paths, when creating the partition locks the software also shall include a locking hierarchy as is known to those skilled in the art to prevent deadlock scenarios. In an illustrative embodiment, when acquiring multiple locks, the lowest numbered 5 partition lock is acquired first and other locks are acquired in order through the highest numbered partition lock and the global lock is acquired last. A higher numbered partition lock, or the global lock, can be acquired when a lower numbered partition lock is held. If a lower numbered partition lock than one that is already held must be acquired, then the higher numbered partition locks are released and reacquired in the manner described above. The locks can be released in any order. As indicated above, the concept of a locking hierarchy is well understood in the field of computer science, see for example, J.W. Havender, Avoiding Deadlock in Multitasking Systems, IBM Systems Journal 7,2 (1968), 74-84

Referring now to FIG. 6, there is shown the code path from FIG. 3 but which is modified to include acquiring and releasing of the N partition locks according to the established locking hierarchy. As provided in the above-described locking hierarchy, all of the partition locks are acquired, STEP 502, starting with partition #1 and continuing until reaching the Nth partition. Thereafter, the global lock is acquired, STEP 504.

20 After acquiring all of the partition locks and the global lock, the software includes code for accessing the data in the locked data set, STEP 506. When all of the partition locks and the global lock are acquired, all of the data within the data set is protected such that it does not change while the data within the data set is being accessed. Once data accessing is completed, all of the partition locks are released,

STEP 508 and the global lock is released, STEP 510. The modification of the software to acquire all-locks for accessing locked data in the system memory, maintains the software code in an operational state, functionally equivalent to the original state of the unaltered software. Such an operational state of the software is maintained even 5 in cases where other activities or actions are taken in accordance with the teachings of the present invention to optimize locking requirement in selected ones of the plurality or more of code paths.

According to this aspect of the present invention, the software is written/revised so that at least selected ones of the plurality or more code paths thereof are optimized with regards to the locking requirements of these selected code paths so as to reduce the contention for acquiring the locks for accessing data within a data set. Thus, each of the various code paths that are accessing data from a given data set are evaluated to determine the code path which experiences the heaviest use, hereinafter the hottest code path, STEP 408. Such a determination is effected using any of a number of methods or techniques known to those skilled in the art, including, but not limited to instrumentation of the software, such as by the addition of counters; or analysis with commercially available tools such as VTUNE © Intel.

The data of the data set that ordinarily would be completely locked for the identified code path is then evaluated to determine the subset of this locked data set 20 that would be accessed within the identified code path. After identifying this data all or some of the identified subset of the locked data is then partitioned or in effect moved into a partition, STEP 410. Reduction of lock contention according to this aspect of the present invention involves moving items from the initial set to either one of the partition sets, or the global set. A data item is moved by assigning it to a

partition, or by determining that no partition makes sense and moving it to the global set. Assigning an item to a partition does not necessarily require a software change because it is simply redefining the locking requirements for that data item. An item placed in a partition may either have an explicit partition number, or be implicitly understood by the software to be in a particular partition. This partition information must be available to the software so it can determine what partition to lock when the item is accessed. When an item is an instantiation of an object, a partition field may be added to the object to explicitly indicate the partition the item resides in. There is shown in FIG. 5C an intermediate state of the data set where some items have been partitioned (i.e., moved to a partition) and others still exist in the initial set. There also is shown in FIG. 5C a data item (i.e., 2____) that is being moved to the global set. In this way, the locking requirements for this item can be reduced from the all-locks configuration (i.e., all partition and global locks) to just the global lock. Because data has been moved to one or more partitions, the global lock does not affect all data of the original data set.

Thus, for the above phonebook example, the data comprising the data set is partitioned in 26 partitions, where the boundaries of each partition is established based on the first letter of the last name of each subscriber. In the case of the above memory management system example, when a software path needs to modify a page descriptor it must be able to determine the partition of the page descriptor so it knows which partition to lock. The page descriptor object structure can be modified to contain a partition field defining the partition that the page descriptor belongs to, or the partition of the page descriptor can be derived from the location of the page within the memory range. Suppose that the global lock in this example also protects a count

5
10
15

20

of the total pages described by these page descriptors. Since the count doesn't conform to any of the partitions, it can be moved to the global set so that only the global lock need be held to access it. Thus, when a path adds a page descriptor it can acquire the global lock, increment the count, and release the global lock while holding
5 the partition lock of the new page descriptor. Alternatively, it may be desirable to create N page count objects where each contains the count of the page descriptors in a partition. These N page count objects can then be moved to their corresponding partitions eliminating the need to acquire the global lock when updating the per-partition page count.

When an item is moved from the initial set to a partition set (i.e., by defining a partition for the item), the locking requirements in the at least selected ones of the software code paths can be reduced. Thus, following data partitioning, the locking requirements for the identified code path are optimized so only the locks associated with the subset of data required for this code path is locked, STEP 412. In other words, the locking requirements of the identified code path is modified from acquiring all locks to acquiring only the locks (i.e., partition locks and/or global lock) needed to lock the data items (i.e., data subsets or partitions) to be accessed in the identified code path. For purposes of the present invention, a software code path according to the present invention, begins at a location in the software code that acquires the
20 global lock (i.e., before the locking is modified), and ends when the global lock is released (i.e., before the locking is modified). The code path can branch and loop, and have multiple release points. The code paths also can share code (i.e., common routines) with other code paths.

As indicated above, after partitioning the data set amongst the N partitions and the global lock set and after optimizing the locking requirements of a code path, the software remains functional or operational. Thus, after optimizing the locking requirements of the identified code path, the programmer(s) or software designer(s) 5 also can perform any or a number of tests on the so-modified software to determine the overall operational performance of the modified software. Such tests can be any of a number of tests known to those skilled in the art, any of a number of tests known in the art that can be adapted for use by any one skilled in the art, or can be developed by one of ordinary skill in the art. The programmer(s) or software designer(s) also evaluate the operational performance of the software to determine if the overall performance of the software is adequate for the intended use, STEPS 10 414,416. In the case where an existing, working piece of code is being modified or upgraded, the so-modified code is evaluated to determine if the overall operational performance has been improved from the earlier version of the code as intended and/or meets any improvement goals that may have been established, expected or intended. Such testing and evaluation of the operational performance is achievable, because the code path(s) of the code comprising the software that access data which were not optimized, remain in operational condition as a result of the “all locks” code modification of STEP 406. In other words, data can be partitioned and the locking 15 requirements of selected code paths can be optimized because of such partitioning without effecting the operational capabilities of the other non-optimized code paths. Furthermore, this provides a mechanism by which the incremental improvement in overall performance of the software resulting from the optimization of by code path locking requirements can be determined and evaluated. Such evaluation of 20

incremental performance also provides a mechanism to perform a cost-benefit analysis to determine if the incremental increase in performance warrants the time taken to optimize locking requirements for a working piece of software.

As indicated above, the programmer(s) or software designer(s) evaluate the overall performance of the modified software (i.e., applications programs/ operating system) so as to determine if the increase in performance is acceptable, STEP 416. If the overall increase is determined to be acceptable (YES, STEP 416) then the programmer(s) or software designers(s) end the process of optimizing the software to reduce the contention of a highly contended lock(s), STEP 420.

If the overall increase is determined not to be acceptable (NO, STEP 416) then the next heaviest code path (i.e., the next hottest code path) is identified, STEP 418. Thereafter, STEPS 410-416 are repeated for the next heaviest code path. This process is repeated until the overall performance of the software is determined to be acceptable (YES, STEP 416) or until the optimization process has exhausted all code paths accessing data of the date set. Thereafter, the programmer(s) or software designers(s) end the process of optimizing the software to reduce the contention of a highly contended lock(s), STEP 420.

It is not generally necessary to identify and remove all items from the initial set to achieve reduced contention and improved performance. In order to determine the most advantageous items to partition, the most often used paths need to be identified. If these paths have multiple branches, the most commonly taken branches must be identified. The items used by the most commonly called paths and most often taken branches are the items that when partitioned will typically result in the greatest reduction in contention. Thus, when a data item is partitioned it is not necessary to

adjust the locking of all paths using or accessing that data item. The paths that are rarely called may continue to lock all partitions because they will have little or no effect on overall contention or overall system performance.

When the locking requirements for a heavily used code path is reduced by partitioning the items it touches, the new locking requirements must be determined and the path modified to only acquire the necessary locks. The correct locking may be any combination of partition locks with or without the global lock, or it may be all locks. The code path determines the correct locking by reading the locales of the items it will touch, or implicitly knowing the partitions of the items it needs locked. The path may or may not be able to determine ahead of time what branch will be taken. It may be possible to lock a known set of needed partitions and wait until later in the path to adjust the locking as needed consistent with the established locking hierarchy. If the items in the path are used in such a way that this is not appropriate then the lock must be acquired at the beginning of the path even if it is not required by the branch that is taken.

In order to reduce the number of partition locks that must be acquired by a particular code path, all the data items used by that code path should be in the same partition. To accomplish this, data items that have an explicit partition number can be dynamically moved from one partition to another at run time. Generally it is necessary to hold the partition lock of both the old partition and the new partition when moving an item. When acquiring the partition lock of a data item that can move between partitions, the partition of the item must be rechecked after the lock was acquired to insure it did not move while it was being locked.

Through this process of moving items out of the initial set and adjusting the locking of the important paths, the so-modified/ updated software becomes more distributed and contention for the global lock is reduced. It is useful to note that after each step of partitioning items and optimizing the locking, the software remains in working condition and the performance and lock contention can be measured.

There is shown in FIG. 7 a high-level flow diagram that illustrates another process or method of the present invention. This method differs from that of FIG. 5 in that the locking requirements for all code paths accessing the data set stored in system memory are optimized so the software is ideally in a state of optimal locking. For the details of STEPS 702, 704, and 712 reference shall be made to the foregoing discussion regarding STEPS 402, 404 and STEP 420 of FIG. 4.

Following the determination of the partitioning methodology (STEP 702), and the creation of the N partition locks (STEP 704) the process according to the alternative embodiment includes partitioning some of the data set, STEP 706. This also changes the locking requirements for all code paths that touch the newly partitioned data. Reference also shall be made to the foregoing discussion for STEP 410 as to the other details for the partitioning of data.

After partitioning of the data, the locking requirements for all code paths that touch the partitioned data are modified, STEP 708. In other words, each code path is examined and, if necessary, modified to match the new locking requirements.

Reference also shall be made to the foregoing discussion regarding STEP 412 as to other details for the modification of the locking requirements of a code path.

After data is partitioned and code paths are changed, the change in overall performance of the software is evaluated to determine if the overall performance of the

so-modified software is acceptable or has reached a desired level, STEP 710. If the overall performance of the software is acceptable or reached the desired level (YES, STEP 710) then the programmer(s) or software designer(s) end the process of optimizing the software to reduce contention, STEP 712. Reference also shall be made to the foregoing discussion regarding STEPS 414, 416 as to other details concerning the evaluation and determination of the modified code's acceptability.

If the overall performance is not acceptable or has not reached the desired level (NO, STEP 710), then STEPS 706-708 are repeated until the overall performance of the modified code is determined to be acceptable (YES, STEP 710) or until all of the data has been partitioned. This alternative embodiment has the beneficial effect that after each iteration of partitioning data and modifying code path locking, the locking of all code paths is optimal given the current partitioning of data. In comparison to the methodology illustrated in FIG. 4, in the alternative embodiment, the programmer or software designer does not have the option to only modify the locking requirements of the heaviest or selected code path. Instead, in the alternative embodiment, all code paths are examined on each iteration to insure that the software remains in a working state.

If the software code paths access more than one data set, then either of the foregoing processes is repeated for each data set being accessed.

Although a preferred embodiment of the invention has been described using specific terms, such description is for illustrative purposes only, and it is to be understood that changes and variations may be made without departing from the spirit or scope of the following claims.

What is claimed is:

1. A method for reducing contention of a highly contended software lock protecting data items of a data set, all of the data items being stored in a system memory of a multi-processor computer system, said method comprising the steps of:
 - creating N partition locks, one for each partition, where $N \geq 2$;
 - identifying one code path from one or more code paths of a software program that access one or more of the data items;
 - determining which data items of the data set are touched by the identified code path;
 - partitioning at least some of the data items that are touched by the identified code path; and
 - optimizing the locking requirements of the identified code path so the locks being acquired and released in the identified code path are those associated with the data items being touched by the identified code path.
2. The method according to claim 1, further comprising the step of modifying the locking requirements of the one or more code paths of the software program that access one or more of the data items so as to acquire all N partition locks and the global lock where a global lock would have been acquired prior to accessing of the one or more data items and so as to release all N partition locks and the global lock where a global lock would have been released after accessing of the one or more data items.

3. The method according to claim 1, wherein the identified code path includes a plurality of branches, and wherein said optimizing includes optimizing the locking requirements of the identified code path so the locks being acquired and released in the code path are those associated with the data items being touched by each branch of identified code path.

4. The method according to claim 3, wherein said optimizing includes optimizing the locking requirements of each branch of the identified code path so the locks being acquired and released in each branch are those associated with the data items being touched by said each branch.

5. The method according to claim 1, further comprising the step of evaluating the software program after said optimizing the locking requirements so as to determine if the overall performance of the software program is acceptable.

6. The method according to claim 5, wherein in the case where said evaluating determines that the overall performance of the software program is not acceptable, then said method includes identifying another code path of the one or more code paths and repeating said steps of determining, partitioning, optimizing, and evaluating for the another identified code path.

7. The method according to claim 1, further comprising the step of first determining a methodology for partitioning the data set.

8. The method according to claim 1, wherein the code path first identified is the heaviest used code path.

9. The method according to claim 6, wherein the code path first identified is the heaviest used code path and wherein the another code path and subsequent code paths are identified sequentially in the direction from the heaviest used code path to a lesser used path.

10. The method according to claim 1, wherein there is one of a plurality or a multiplicity of code paths that access one or more of the data items.

11. A method for reducing contention of a highly contended software lock protecting data items of a data set, all of the data items being stored in a system memory of a multi-processor computer system, said method comprising the steps of:
first determining a methodology for partitioning the data set;
creating N partition locks, one for each partition, where $N \geq 2$;
modifying the locking requirements of each of one or more code paths of a software program that access one or more of the data items so as to acquire all N partition locks and a global lock where the global lock would have been acquired prior to accessing of the one or more data items and so as to release all N partition locks and the global lock where the global lock would have been released after accessing of the one or more data items;

identifying one code path from the one or more code paths of the software program that access one or more of the data items; next determining which data items of the data set are touched by the identified code path; partitioning at least some of the data items that are touched by the identified code path; and

optimizing the locking requirements of the identified code path so the locks being acquired and released in the identified code path are those associated with the data items being touched by the identified code path.

12. The method according to claim 11, wherein the identified code path includes a plurality of branches, and wherein said optimizing includes optimizing the locking requirements of the identified code path so the locks being acquired and released in the code path are those associated with the data items being touched by each branch of identified code path.

13. The method according to claim 12, wherein said optimizing includes optimizing the locking requirements of each branch of the identified code path so the locks being acquired and released in each branch are those associated with the data items being touched by said each branch.

14. The method according to claim 11, further comprising the step of evaluating the software program after said optimizing the locking requirements so as to determine if the overall performance of the software program is acceptable.

15. The method according to claim 14, wherein in the case where said evaluating determines that the overall performance of the software program is not acceptable, then said method includes identifying another code path of the one or more code paths and repeating said steps of determining, partitioning, optimizing and evaluating for the another identified code path.

16. The method according to claim 15, wherein the code path first identified is the heaviest used code path and wherein the another code path and subsequent code paths are identified sequentially in the direction from the heaviest used code path to a lesser used path.

17. A method for reducing contention of a highly contended software lock protecting data items of a data set, all of the data items being stored in a system memory of a multi-processor computer system, said method comprising the steps of:
creating N partition locks, one for each partition, where $N \geq 2$;
partitioning some of the data items; and
modifying the locking requirements of all code paths of the one or more code paths of a software program that access one or more of the data items so that the locks being acquired and released in each of said all code paths are those associated with the touching data items.

18. The method according to claim 17, wherein one code path of said all code paths includes a plurality of branches, and wherein said modifying includes modifying the locking requirements of each branch of said one code path so the locks being acquired and released in each branch are those associated with the data touching said each branch.

19. The method according to claim 17, further comprising the step of evaluating the software program after said modifying the locking requirements so as to determine if the overall performance of the software program is acceptable.

20. The method according to claim 19, wherein in the case where said evaluating determines that the overall performance of the software program is not acceptable, then said method includes partitioning more data items and repeating said steps of modifying and evaluating.

21. The method according to claim 17, further comprising the step of first determining a methodology for partitioning the data set.

22. The method according to claim 17, wherein there is one of a plurality or a multiplicity of code paths that access one or more of the data items.

23. A method for reducing contention of a highly contended software lock protecting data items of a data set, all of the data items being stored in a system memory of a multi-processor computer system, said method comprising the steps of:

- first determining a methodology for partitioning the data set
- creating N partition locks, one for each partition, where $N \geq 2$;
- partitioning some of the data items;
- modifying the locking requirements of all code paths of the one or more code paths of a software program that access one or more of the data items so that the locks being acquired and released in each of said all code paths are those associated with the touching data items; and
- evaluating the software program after said modifying the locking requirements so as to determine if the overall performance of the software program is acceptable.

24. The method according to claim 23, wherein in the case where said evaluating determines that the overall performance of the software program is not acceptable, then said method includes partitioning more data items and repeating said steps of modifying and evaluating.

25. A method for accessing one or more data items in a data set by a software program, all of the data items being stored in a system memory of a multi-processor computer system, said method comprising the steps of:

creating N partition locks, one for each partition, where $N \geq 2$;

identifying one code path from one or more code paths of the software program that access one or more of the data items;

determining which data items of the data set are touched by the identified code path;

partitioning at least some of the data items that are touched by the identified code path;

optimizing the locking requirements of the identified code path so the locks being acquired and released in the identified code path are those associated with the data items being touched by the identified code path;

locking the data items of the data set that are touched by the identified code path while keeping unlocked the data items of the data set that are not being touched by the identified code path;

accessing one or more of the locked data items; and

releasing the locks associated with the locked data.

26. The method according to claim 25, further comprising the step of modifying the locking requirements of the one or more code paths of the software program that access one or more of the data items so as to acquire all N partition locks and the global lock where a global lock would have been acquired prior to

accessing of the one or more data items and so as to release all N partition locks and the global lock where a global lock would have been released after accessing of the one or more data items.

27. The method according to claim 25, wherein in the case where accessing data in the one or more code paths in which all N partition locks and the global lock are acquired and released, said method further comprised the steps of:

- acquiring all N partition locks and the global lock;
- accessing the data being protected by the acquired N partition locks and the global lock; and
- releasing all N partition locks and the global lock.

28. The method according to claim 25, wherein the identified code path is the heaviest used code path that accesses data items.

29. The method according to claim 25, wherein locking requirements for a plurality of code paths are optimized, and wherein said acquiring, accessing and releasing are selectively effected in any one of the plurality of code paths provided that the data items to be locked in said any one code path are not locked in any other of the plurality of code paths.

ABSTRACT

Featured is a method or process for reducing contention of a highly contended software lock(s) that is protecting multiple data items, where the software has a plurality of code paths accessing the data items. The method includes creating additional partition locks to protect subsets of the data items protected by the existing global lock. Such a method further includes acquiring all partition locks and the global lock, wherever a global lock would have been acquired to protect data. The method also includes identifying one or more heavily used code paths and determining which data items are touched by the identified one or more heavily used code paths. These data items are then moved into a partition, if they were not partitioned earlier. The locking requirements for each of the identified one or more heavily used code paths are optimized to match the reduced locking requirements because of the partitioned data items. In other words the locking requirements are reduced so only the locks for the partitions including the data items that touch the code path are acquired. In more specific embodiments, the so-modified software is evaluated to determine if there is an acceptable increase in overall system performance resulting from the optimization of the locking requirements. If the modified system's performance is not acceptable, then the next most heavily used code path is identified and the locking requirements for this code path are optimized. Such optimization is continued until the system exhibits an acceptable overall system performance or all code paths are optimized.





FIG. I
PRIOR ART

FIG. 2
PRIOR ART







FIG. 3
PRIOR ART




FIG. 4

Fig 5A




Fig 5B





Fig 5C






FIG. 6




Fig 7

DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I hereby declare that: My residence, post office address and citizenship are as stated below next to my name. I believe I am the original, first and sole inventor (if only one name is listed at 201) below or an original, first and joint inventor (if plural names are listed at 201-208 below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

**METHOD OF REDUCING CONTENTION OF A HIGHLY CONTESTED LOCK PROTECTING
MULTIPLE DATA ITEMS**

which is described and claimed in:

- the specification attached hereto.
- the specification in U.S. Application Serial Number _____, filed on _____.
- the specification in PCT international application Number _____, filed on _____; and was amended on _____.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a). I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed.

Prior Foreign/PCT Applications and Any Priority Claims Under 35 U.S.C. §119:

Application No.	Filing Date	Country	Priority Claimed Under 35 U.S.C. §119?
			<input type="checkbox"/> YES <input type="checkbox"/> NO
			<input type="checkbox"/> YES <input type="checkbox"/> NO
			<input type="checkbox"/> YES <input type="checkbox"/> NO
			<input type="checkbox"/> YES <input type="checkbox"/> NO
			<input type="checkbox"/> YES <input type="checkbox"/> NO

I hereby claim the benefit under 35 U.S.C. §120 of any United States application(s) or PCT international application(s) designating the United States of America that is/are listed below, and, insofar as the subject matter of each of the claims of this application is not disclosed in that/those prior application(s) in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose material information as defined in 37 CFR §1.56(a) which occurred between the filing date of the prior application(s) and the national or PCT international filing date of this application:

**Prior U.S. Applications or PCT International Applications Designating the U.S-Benefit
Under 35 U.S.C. §120**

U.S. Applications		Status (Check One)		
Application Serial No.	U.S. Filing Date	Patented	Pending	Abandoned
PCT Applications Designating the U.S.				
Application No.	Filing Date	U.S. Serial No. Assigned		

**CLAIM FOR BENEFIT OF PRIOR U.S. PROVISIONAL APPLICATION(S)
(35 U.S.C. §119(e))**

I hereby claim the benefit under Title 35, United States Code, §119(e) of any United States provisional application(s) listed below:

Applicant	Provisional Application Number	Filing Date

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) with full powers of association, substitution and revocation to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Sewall P. Bronstein (Reg. No. 16,919)
David G. Conlin (Reg. No. 27,026)
George W. Neuner (Reg. No. 26,964)
Linda M. Buckley (Reg. No. 31,003)
Peter J. Manus (Reg. No. 26,766)

Peter F. Corless (Reg. No. 33,860)
Cara Z. Lowen (Reg. No. 38,227)
William J. Daley, Jr. (Reg. No. 34,587)
Robert L. Buchanan (Reg. No. 40,927)
David Tucker (Reg. No. 27,840)

Christine C. O'Day (Reg. No. 38,256)
Lisa Swiszcz Hazzard (Reg. No. 44,368)
George Hartnell (Reg. No. 44,368)
Jennifer Holmes (Reg. No. P-46,778)
Kerri Pollard Schray (Reg. No. P-47,066)

SEND CORRESPONDENCE TO:

William J. Daley, Jr.
Dike, Bronstein, Roberts & Cushman
Intellectual Property Practice Group
EDWARDS & ANGELL, LLP
130 Water Street
Boston, Massachusetts 02109

DIRECT TELEPHONE CALLS TO:

(617) 523-3400

2 0 1	FULL NAME OF INVENTOR	LAST NAME	FIRST NAME	MIDDLE NAME
	RESIDENCE & CITIZENSHIP	CITY	STATE OR FOREIGN COUNTRY	COUNTRY OF CITIZENSHIP
	POST OFFICE ADDRESS	POST OFFICE ADDRESS	CITY	STATE OR COUNTRY AND ZIP CODE
	Martin	Brian	James	
	Redmond	Washington	Canada	
	21701 NE 73rd Place	Redmond	Washington 98053	

2 0 2	FULL NAME OF INVENTOR	LAST NAME	FIRST NAME	MIDDLE NAME
	RESIDENCE & CITIZENSHIP	CITY	STATE OR FOREIGN COUNTRY	COUNTRY OF CITIZENSHIP
	POST OFFICE ADDRESS	POST OFFICE ADDRESS	CITY	STATE OR COUNTRY AND ZIP CODE
	Ryan	Michael		
	Redmond	Washington	United States	
	15606 NE 40th Street, Apt. Q362	Redmond	Washington 98052	

I hereby further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Signature of Inventor 201	Signature of Inventor 202
Date:	Date: